

# Distillation protocols for Fourier states in quantum computing

Cody Jones\*

*Edward L. Ginzton Laboratory, Stanford University, Stanford,  
California 94305-4088, USA*

March 14, 2013

## Abstract

Fourier states are multi-qubit registers that facilitate phase rotations in fault-tolerant quantum computing. We propose distillation protocols for constructing the fundamental,  $n$ -qubit Fourier state with error  $O(2^{-n})$  at a cost of  $O(n \log n)$  Toffoli gates and Clifford gates, or any arbitrary Fourier state using  $O(n^2)$  gates. We analyze these protocols with methods from digital signal processing. These results suggest that phase kickback, which uses Fourier states, could be the current lowest-overhead method for generating arbitrary phase rotations.

## 1 Introduction

Fault-tolerant quantum computing seeks to develop quantum information processors that are resilient to faults in any hardware component using quantum error-correction codes [1, 2]. Recent attention has focused on how to minimize the resource costs of essential quantum computing primitives, such as a “universal” set of operations. Theoretical analysis has shown that there must be at least one operation required for universality but not natively supported by the quantum code (often called “non-transversal”) [3, 4]. Under realistic device parameters, the resource costs of the creating non-transversal gates dominate the total overhead for achieving fault tolerance [5–9]. Therefore, choosing the appropriate non-transversal operations directly impacts resource costs for fault-tolerant quantum computers.

Many investigations, including some involving this author, adopt a “simpler is better” approach to selecting which non-transversal gate completes the universal set, so they focus on the single-qubit gate  $T = \exp[i\pi(I - \sigma^z)/8]$ , where  $I$  is identity [5–21]. Recently, Landahl and Cesare proposed using a family

---

\*Electronic address: ncodyjones@gmail.com

of rotation gates that they label  $Z_k = \exp[i\pi\phi(I - \sigma^z)/2^{k+1}]$ , which includes  $Z_2 = T$ , in what they term “complex instruction set computing (CISC)” [22]; a similar approach was used as a subroutine in Ref. [5]. Instead of the  $T$  gate, we promote an alternative, the three-qubit Toffoli gate [2, 23–25]. As an aside, our results also allow one to implement the CISC proposal efficiently. Ultimately, the universal gate set is used to approximate quantum circuits needed in an algorithm. Toffoli gates are already preferred for arithmetic circuits [6, 26–28], but this work goes further to show that Toffoli is also efficient for arbitrary gates.

Recent work has shown that any arbitrary gate can be efficiently approximated using an instruction set that includes the  $T$  gate [12, 20]. It has also been shown that arbitrary gates can be efficiently approximated with phase kickback, which uses Toffoli gates, so long as one has a multi-qubit resource that we call a “Fourier state” [14]. This paper completes the phase kickback method by showing how to construct the Fourier state efficiently. Moreover, we go further to argue that phase kickback can have lower resource costs than methods that use  $T$  gates; when this is true, one should consider using Toffoli as the non-transversal operation in fault-tolerant quantum computing.

To give some context, a Fourier state of size  $n$  qubits is defined as

$$|\gamma^{(k)}\rangle = \frac{1}{\sqrt{N}} \sum_{y=0}^{N-1} e^{i2\pi ky/N} |y\rangle, \quad (1)$$

where  $N = 2^n$ . Note that this is the sign convention of Ref. [2] and opposite of that in Ref. [14]. These states are eigenstates of the modular addition operator  $U_{\oplus 1}|x\rangle = |x + 1 \pmod{N}\rangle$ :

$$U_{\oplus 1}|\gamma^{(k)}\rangle = e^{-i2\pi k/N} |\gamma^{(k)}\rangle. \quad (2)$$

Using this property, a phase-rotation gate can be approximated using a Fourier state and an addition circuit, which is known as phase kickback [14, 29, 30]. This method can produce any rotation around the  $\sigma^z$  axis of the Bloch sphere in units of  $\pi/2^{n-1}$  radians, so the precision required by a quantum algorithm determines  $n$ . A notable feature of phase kickback is that the register  $|\gamma^{(k)}\rangle$  is preserved, which means it can be used repeatedly. Many implementations of addition circuits are known [26–28], but the fault-tolerant preparation of Fourier states has received less attention. Kitaev *et al.* propose a scheme based on phase estimation [30], but this protocol suffers from two notable disadvantages. First, the resulting state  $|\gamma^{(k)}\rangle$  has random odd  $k$ . Second, the protocol requires a Fourier transform, which requires phase rotations; since the purpose of phase kickback is to produce phase rotations, implementing the Fourier transform to produce the Fourier state requires an approximate, iterative procedure. This paper develops a fault-tolerant distillation protocol for producing the frequently used  $|\gamma^{(1)}\rangle$  state with  $O(n \log n)$  gates from a finite set. We also give a related protocol for constructing any  $|\gamma^{(k)}\rangle$  using  $O(n^2)$  gates.

The circuit for constructing Fourier states is implemented in a fault-tolerant quantum computer that, owing to the constraints of error correction [1, 3, 4],

has a limited set of gates. In what follows, we will denote Pauli operators by  $Z \equiv \sigma^z$ , *etc.* Some gates require more resource overhead to produce than others. Gates in the Clifford group are generated by combinations of: Hadamard  $H = (1/\sqrt{2})(X + Z)$ , phase gate  $S = \exp[i\pi(I - Z)/4]$ , and CNOT, up to global phase that we ignore. In the set of “Clifford gates,” we also include ancilla qubits initialized to  $|0\rangle$  and measurement  $M_z$  performed in the  $Z$  basis. We assume that any Clifford gate is economical in terms of resource cost and that non-Clifford gates (those outside the Clifford group) are the dominant cost. This assumption is justified by analysis showing that non-Clifford gates are substantially more resource intensive than Clifford gates [5–9]. Nevertheless, a universal set of operations requires at least one non-Clifford gate, and we select Toffoli. Recent work shows that low-overhead constructions for the Toffoli gate exist, often by implementing low-fidelity  $T$  gates with subsequent error correction [24, 25].

The recent analysis of resource costs in Refs. [7–9, 24] indicates that producing a Toffoli gate with error  $10^{-12}$  requires comparable physical resources when using surface code error correction to producing a single  $T$  gate having the same error probability. Although Toffoli gates could be the more resource-efficient choice of non-Clifford gate, constructing quantum algorithms efficiently with Toffoli gates was not fully resolved (prior to this work) in situations where arbitrary phase rotations are required. A major consequence of this paper is that it completes the phase-kickback protocol by showing that the discrete set of Clifford operations and Toffoli gates can efficiently approximate any quantum unitary, which could replace constructions using  $T$  gates [11, 12, 16, 20].

The paper is organized as follows. Section 2 presents a protocol for distilling the “fundamental” Fourier state  $|\gamma^{(1)}\rangle$  from approximations produced using only Clifford gates. Section 3 analyzes the resource costs of the distillation protocol, which we summarize here. Constructing an  $n$ -qubit  $|\gamma^{(1)}\rangle$  state requires circuit width  $2n + O(1)$  qubits, circuit depth  $O(n)$  gates, and  $O(n \log n)$  Toffoli gates in total. Section 4 compares phase kickback to a competing method using  $T$  gates, which is relevant since the purpose of Fourier distillation is to complete fault-tolerant constructions for phase kickback. Section 5 outlines a protocol for distilling  $|\gamma^{(k)}\rangle$  with arbitrary  $k$  using  $O(n^2)$  Toffoli gates. The paper concludes with a discussion of why the combination of these results and recent work in constructing Toffoli gates makes phase kickback a compelling approach to approximating arbitrary quantum gates with low overhead.

## 2 Distilling the fundamental Fourier state

The fundamental  $|\gamma^{(1)}\rangle$  state of size  $n$  qubits is required for phase-kickback constructions for both single-qubit phase rotations and two-qubit, controlled phase rotations [14]. This state is also useful for constructing a quantum Fourier transform (or its approximate version) through a special form of phase kickback called quantum-variable rotation [14]. This section shows how to construct  $|\gamma^{(1)}\rangle$  using a distillation protocol based on addition circuits. We generalize the method in a later section to distill  $|\gamma^{(k)}\rangle$  for arbitrary  $k$ , but the special case of

$k = 1$  requires fewer quantum gates.

## 2.1 Approximate fundamental Fourier state

Any pure Fourier-basis state where  $N$  in Eqn. (1) is a power of 2 is separable into individual qubits. Using a general  $Z$ -axis rotation  $R_z(\phi) = \exp[i\pi\phi(I - Z)/2]$ , a Fourier state can be decomposed as

$$|\gamma^{(k)}\rangle = [R_z(k\pi/2^0)|+\rangle] \otimes [R_z(k\pi/2^1)|+\rangle] \otimes [R_z(k\pi/2^2)|+\rangle] \otimes (\dots). \quad (3)$$

The single-qubit state  $|+\rangle = H|0\rangle = (1/\sqrt{2})(|0\rangle + |1\rangle)$ . Using Eqn. (3), we can see that the quantum state

$$|\tilde{\gamma}^{(1)}\rangle = Z|+\rangle \otimes S|+\rangle \otimes I|+\rangle \otimes I|+\rangle \otimes (\dots) \quad (4)$$

is an approximation of  $|\gamma^{(1)}\rangle$  (denoted with tilde). Moreover,  $|\tilde{\gamma}^{(1)}\rangle$  can be produced using only Clifford gates. The rotations of  $\pi/4$ ,  $\pi/8$ , *etc.* in Eqn. (3) that are omitted in Eqn. (4) become exponentially close to identity (in gate fidelity) with increasing qubit index, so we approximate them with identity gates. This is the same justification behind neglecting small-angle rotations in the approximate quantum Fourier transform [31]. The fidelity between the approximate and ideal states is  $|\langle \tilde{\gamma}^{(1)} | \gamma^{(1)} \rangle|^2 \geq 0.81$  for all values of  $n$ .

Each approximate state can be expanded in the orthonormal Fourier-state basis as

$$|\tilde{\gamma}^{(1)}\rangle = \sum_{j=0}^{N-1} a_j |\gamma^{(j)}\rangle, \quad (5)$$

where the dominant term among the complex coefficients is  $a_1$ , with magnitude  $|a_1|^2 \approx 0.81$  from above. We ignore complex phase because we always work in the Fourier basis, and our distillation protocol depends only on the magnitudes of the Fourier-basis coefficients.

## 2.2 Distillation protocol

Using two approximate  $|\tilde{\gamma}^{(1)}\rangle$  states, the distillation protocol is very simple. First, add one register to the other. Binary-encoded, mod- $2^n$  addition given by

$$U_{\text{add}}|v\rangle|w\rangle = |v\rangle|w + v \pmod{2^n}\rangle \quad (6)$$

has been studied extensively [26–28]. Notably, many addition circuits use the Toffoli gate as the non-Clifford operation. In the Fourier basis, the addition circuit implements

$$U_{\text{add}}|\gamma^{(k)}\rangle|\gamma^{(k')}\rangle = |\gamma^{(k-k')}\rangle|\gamma^{(k')}\rangle. \quad (7)$$

As an aside, this is precisely phase kickback, where the Fourier index  $k'$  of the second register determines the quantum-variable rotation applied to the

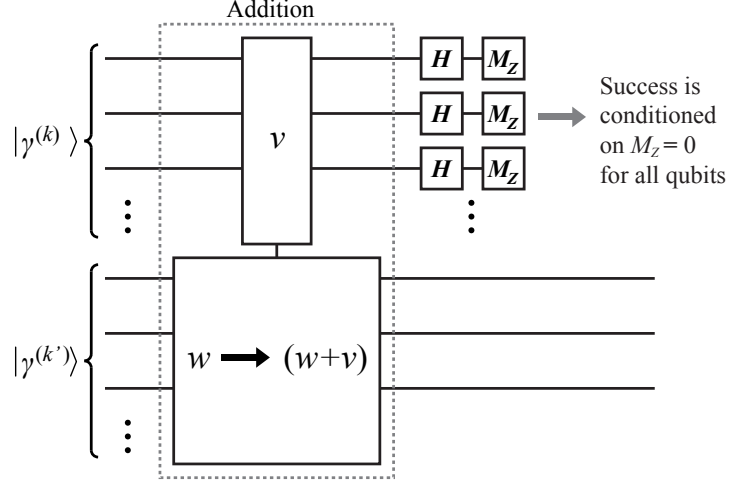


Figure 1: Circuit for distilling Fourier states. Each Fourier state is a multi-qubit register, and only three qubits are shown, with the rest indicated by vertical dots. For clarity, this circuit shows the inputs as pure Fourier-basis states, but in actual distillation protocols, the inputs will be mixed states. The addition circuit shown in the dashed box would be decomposed into Clifford gates and Toffoli gates. The verification step is implemented with Hadamard  $H$  and computational-basis measurement  $M_z$  on each qubit in the top register.

first register (see Section 4.1 of Ref. [14]). Second, measure the first register in the Fourier basis, and postselect the cases where the result is  $|\gamma^{(0)}\rangle$ . The resulting output has each of its Fourier-basis coefficients  $\{a_k\}$  weighted by the probability that *both* inputs to distillation were in the  $|\gamma^{(k)}\rangle$  state. If both inputs had sizable overlap with a particular state, then the fidelity conditioned on successful distillation is concentrated to a higher magnitude, and probability of being in unwanted basis states is suppressed.

The quantitative expressions for distillation success probability and projection of output state into the Fourier basis are simple to derive. Let the two inputs to distillation have Fourier-basis coefficients  $\{a_j\}$  and  $\{a'_j\}$  as in Eqn. (5). The probability of measuring  $|\gamma^{(0)}\rangle$  (*i.e.* distillation succeeds) is given by

$$P_{\text{success}} = \sum_{y=0}^{N-1} |a_y|^2 |a'_y|^2. \quad (8)$$

The output register of distillation will have Fourier coefficients  $\{b_j\}$  with magnitudes

$$|b_j|^2 = \frac{|a_j|^2 |a'_j|^2}{P_{\text{success}}}. \quad (9)$$

These expressions mirror those of entanglement distillation [32,33]. We measure

the fidelity of the output state as  $F = |b_1|^2$ , and the error probability in the distilled state is  $\epsilon = 1 - F$ .

A general Fourier-basis measurement would pose a problem because it requires operations outside the Clifford group, but we show how to circumvent this issue with Clifford gates. Since the quantum computer only supports computational-basis measurements, we would require a quantum Fourier transform (QFT) to map between the bases. This is essentially the obstacle encountered by the Kitaev-Shen-Vyalyi protocol [30], which addresses the matter with an iterative procedure of approximate QFTs. However, our distillation protocol does not require a *complete* Fourier-basis measurement; instead, we only need to know if the first register is in state  $|\gamma^{(0)}\rangle$ . This state happens to be the tensor product of  $|+\rangle$  states, which are eigenstates in the  $X$ -basis. Hence, we only require application of the Hadamard gate  $H$  followed by measurement  $M_z$  on each qubit in the first register in Fig. 1. If each qubit is the  $|+\rangle$  state, then the register was projected into  $|\gamma^{(0)}\rangle$ , and distillation succeeds. Otherwise, reject the output and attempt again. Since  $H$  and  $M_z$  are in the Clifford group, they are considered inexpensive to produce relative to the preceding non-Clifford addition circuit. In addition to preparing the specific state  $|\gamma^{(1)}\rangle$  (as opposed to a random Fourier-basis state [14, 30]), this measurement trick is how our protocol improves on the method in Ref. [30].

If the two input states are both  $|\tilde{\gamma}^{(1)}\rangle$  from Eqn. (4), then  $|a_j|^2 = |a'_j|^2$  for all  $j$ . In general, when the inputs satisfy  $|a_j|^2 = |a'_j|^2$  for all  $j$ , we say the distillation is “symmetric.” The success probability is  $P_{\text{success}} = \sum_{y=0}^{N-1} |a_y|^4 \approx 0.67$ , where the coefficients  $\{a_j\}$  can be calculated using the method in Section 2.3. The Fourier-basis weights of the output state are  $|b_j|^2 = |a_j|^4 / P_{\text{success}}$ . The fidelity  $F = |b_1|^2$  after one round of distillation is upper-bounded by 0.986, so multiple rounds of distillation are needed to reach arbitrarily high fidelity. This bounded fidelity also means that early rounds of distillation can use fewer than  $n$  qubits to represent the intermediate Fourier states, as we explain in Section 3; before explaining that technique, we must quantify the fidelity in each round.

We define an  $n$ -qubit,  $m$ -round distilled  $|\tilde{\gamma}_m^{(1)}\rangle$  Fourier state as having “sufficiently high fidelity” if its fidelity with the pure Fourier state satisfies

$$1 - \left| \langle \tilde{\gamma}_m^{(1)} | \gamma^{(1)} \rangle \right|^2 \leq \sin^2(\pi/2^n). \quad (10)$$

Subscript denotes how many rounds of symmetric distillation have been successfully applied, so initial state  $|\tilde{\gamma}^{(1)}\rangle = |\tilde{\gamma}_0^{(1)}\rangle$ . In phase kickback, the constraint in Eqn. (10) represents the highest accuracy that is needed. When the  $|\tilde{\gamma}_m^{(1)}\rangle$  register is used for phase kickback, there are two sources of error that we consider here. The first is that the register  $|\tilde{\gamma}_m^{(1)}\rangle$  is not pure, meaning it has non-zero overlap with some other Fourier basis state. The second error source is that any phase rotation is truncated to  $n$  bits of precision. As a result, the truncated angle error is at most  $\pi/2^n$  radians, which results in an upper bound on the rotation-gate error probability of  $\sin^2(\pi/2^n) \approx (\pi/2^n)^2$ . In phase kickback

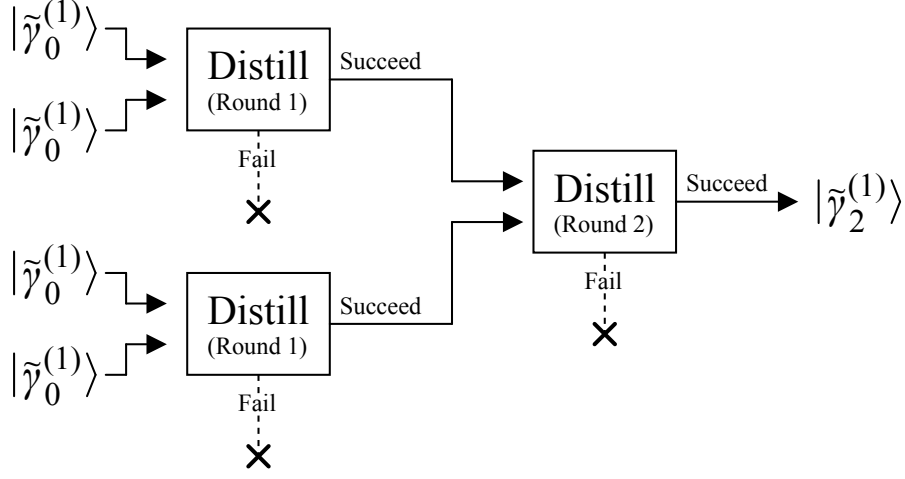


Figure 2: Symmetric distillation of Fourier states, where two rounds are shown. The approximate initial states are defined in Eqn. (4). Each distillation step is implemented using the quantum circuit in Fig. 1. Distillation is probabilistic, and when a distillation circuit fails, that circuit and preceding steps that feed into it must be repeated. The probability of failure decreases super-exponentially in round number, so the overhead of distillation failure is bounded.

using an  $n$ -qubit  $|\tilde{\gamma}_m^{(k)}\rangle$  state, the combination of the two errors means that any resulting rotations are accurate to  $\pm\pi/2^{(n-1)}$  radians, or at least  $(n-1)$  bits. Ultimately,  $n$  is chosen based on the gate-accuracy requirements of the quantum algorithm. We arbitrarily choose to balance the error from a noisy  $|\tilde{\gamma}_m^{(1)}\rangle$  with the worst-case truncation-of-angle error. If  $|\tilde{\gamma}_m^{(1)}\rangle$  is used for other applications, such as complex-instruction-set quantum computing [22], a different accuracy may be required.

The first round of distillation will produce a Fourier state accurate to about 5 bits. To construct an  $n$ -qubit Fourier state, we develop a distillation protocol consisting of multiple rounds of symmetric distillation. The symmetric distillation subroutines are arranged in a binary tree, as shown in Fig. 2. Multiple low-fidelity  $|\tilde{\gamma}_0^{(1)}\rangle$  input states are distilled to arrive at a single output state,  $|\tilde{\gamma}_r^{(1)}\rangle$ . Subscript  $r$  is the number of rounds of symmetric distillation, or the depth of this binary tree arrangement. Choosing  $r$  depends on the desired number of precision qubits  $n$  in the Fourier state. In the next section, we show that  $r$  scales as  $O(\log n)$ .

### 2.3 Fourier analysis and distillation efficiency

The distillation protocol can be understood by viewing probability amplitudes of the input state in the computational basis as discrete samples in time of a

function  $f(t) = e^{i2\pi\nu(t)}$  such that  $|\tilde{\gamma}_0^{(1)}\rangle = \sum_{y=0}^{N-1} f(y/N)|y\rangle$ . In this picture, the probability amplitudes of the quantum state in the Fourier basis are related to Fourier-series coefficients  $\{c_j\}$  given by

$$f(t) = \sum_{k=-\infty}^{\infty} c_k e^{i2\pi kt}. \quad (11)$$

The correspondence exists because a quantum Fourier transform maps between computational and Fourier bases. The number of “samples” is  $N = 2^n$ , the number of computational basis states. Discrete sampling causes aliasing according to the Shannon-Nyquist Theorem, so Fourier-basis probability amplitudes  $\{a_j\}$  are related to the Fourier series of  $f(t)$  by

$$a_j = \sum_{x=-\infty}^{\infty} c_{(Nx+j)}. \quad (12)$$

If  $N$  is sufficiently large (*e.g.*  $n \geq 6$ ), then  $a_j \approx c_j$  for  $0 \leq j < N/2$  or  $a_{(N+j)} \approx c_j$  for  $-N/2 \leq j < 0$ , because the coefficients  $c_j$  decay in magnitude asymptotically as  $|c_j|^2 \propto 1/j^2$ , which means the error from neglecting aliased frequencies is suppressed exponentially in  $n$ . This asymptotic upper bound follows from Parseval’s theorem for any signal that is square-integrable over its period, a condition which corresponds to normalized quantum states.

In each approximate  $|\tilde{\gamma}_0^{(1)}\rangle$  register, the first qubit is the most significant bit in a binary encoding of equally-spaced time coordinates for samples of  $f(t)$ . By using only  $Z$  and  $S$  rotations, we are effectively discretizing the phase of  $f(t) = e^{i2\pi\nu(t)}$  to two bits of precision as a piecewise-constant function over four equally-sized intervals in the domain  $t \in [0, 1)$ . We can readily calculate the  $j^{\text{th}}$  Fourier series coefficient of this function:

$$\begin{aligned} c_j &= \int_0^1 e^{i2\pi(\nu(t)-jt)} dt \\ &= \sum_{r=1}^4 \int_{(r-1)/4}^{r/4} e^{-i2\pi(jt-(r-1)/4)} dt \\ &= \left( \frac{1-i}{2\pi j} \right) \sum_{m=1}^4 e^{-i\pi m(j-1)/2}, \end{aligned} \quad (13)$$

which is valid everywhere except  $j = 0$ , in which case  $c_0 = 0$ . Sign convention follows Eqn. (1). The only nonzero terms occur for  $j \equiv 1 \pmod{4}$ , and they are  $c_j = (2-2i)/(\pi j)$ . The squared magnitudes of the largest Fourier components for state  $|\tilde{\gamma}_0^{(1)}\rangle$  are plotted in the spectrum in Fig. 3. The expression for Fourier-series coefficients allows us to derive bounds on distillation performance and hence the necessary number of rounds of distillation. For example, since there is no relative phase between these coefficients, initial Fourier states have  $|a_1|^2 > |c_1|^2$ , where  $|c_1|^2 = 8/\pi^2$ .






Figure 3: Frequency spectrum of  $f(t)$ . The height of a bar at  $j$  corresponds to the squared magnitude  $|c_j|^2$ , where  $\{c_j\}$  are Fourier-series coefficients of  $f(t)$ . Eqn. (12) relates these series coefficients to Fourier-basis amplitudes.

The output state of the distillation protocol, conditioned on success, has modified Fourier components. In symmetric distillation, the relative magnitude of each component to the fundamental harmonic is squared. After normalization, the largest component at  $j = 1$  is amplified, while the rest are suppressed. The second-largest component is the “sideband” at  $j = -3$  in Fig. 3. Note that because of aliasing in the frequency spectrum,  $c_{(-3)}$  maps to  $a_{(N-3)}$ ; we assume that Fourier-series term  $c_{(N-3)}$  is negligibly small.

Distillation proceeds until the sidebands are suppressed to a sufficiently low level. The rate at which these sidebands are suppressed dictates how many rounds of distillation are required, which determines the total number of gates in the protocol. This rate is limited by the ratio in magnitudes between the fundamental  $j = 1$  harmonic and the second-largest sideband at  $j = -3$ . This behavior is analogous to the rate of convergence in Markov-chain Monte Carlo, which depends on the magnitude of the second-largest eigenvalue of the state transition matrix (the largest eigenvalue of a stochastic matrix is 1). Successful symmetric state distillation through  $r$  rounds modifies each Fourier-basis amplitude from  $a_j$  to  $b_j^{(r)}$  according to

$$|b_j^{(r)}|^2 = \frac{1}{C} (|a_j|^2)^{2^r}, \quad (14)$$

where

$$C = \sum_{j=0}^{N-1} (|a_j|^2)^{2^r} \quad (15)$$

is the normalization. Since any sidebands to the fundamental harmonic ( $j = 1$  in this case) will be suppressed super-exponentially in  $r$ , we need to only focus on the magnitude of the largest sideband at  $c_{(-3)}$ , which will dominate the error in the output of distillation. As a result, the error  $\epsilon = 1 - \left| \langle \tilde{\gamma}_m^{(k)} | \gamma^{(k)} \rangle \right|^2$  in the distilled Fourier state is closely approximated by

$$\epsilon \approx \left( |c_{(-3)}|^2 / |c_1|^2 \right)^{2^r}. \quad (16)$$

Consequently, the ratio  $|c_{(-3)}|^2 / |c_1|^2 = 9$  (exactly) dictates how fast error is suppressed through distillation. Since we require that  $\epsilon \leq (\pi/2^n)^2$ , we can determine the number of rounds of distillation as

$$R = \left\lceil \log_2 \left( \frac{2n - 2 \log_2 \pi}{\log_2(|c_1|^2 / |c_{(-3)}|^2)} \right) \right\rceil. \quad (17)$$

This expression can be simplified to  $R = \lceil \log_2(0.63n - 1.04) \rceil$  (approximately), which shows that  $R$  scales as  $O(\log n)$ . Moreover, Eqn. (16) shows that the error at the output of each successive round of distillation is squared. Eqn. (10) shows that the number of qubits needed to represent an approximate Fourier state is  $O(\log \epsilon)$ , so the smallest allowable size in qubits of intermediate distilled states will double after each round. The next section uses this technique to save resources.

### 3 Resource analysis for distilling the fundamental Fourier state

This section shows that distilling the fundamental  $n$ -qubit Fourier state is efficient in the sense that it requires at most  $O(n \log n)$  Toffoli gates and total gates, with circuit width  $2n + O(1)$  qubits. After each round of distillation, the number of bits of precision in the Fourier states doubles, so we also double the number of qubits going into the next round. The procedure is: (a) after one round of distillation, each Fourier state is accurate to  $s$  qubits; (b) append  $s$  more qubits in the  $|+\rangle$  state to each register; (c) repeat distillation on the input states of size  $2s$  qubits. The additional error of making a larger approximate Fourier state by appending  $s$  qubits in the  $|+\rangle$  state is less than the error already present, so the fidelity is not reduced appreciably. The extra qubits provide space for the output state of distillation (if it succeeds) to contain twice as many accurate qubits.

Each round of distillation on  $s$ -qubit registers uses addition circuits that each require  $(2s - 4)$  Toffoli gates [26] (note that the carry-out Toffoli is unnecessary and removed). If there are  $R$  rounds of distillation, then the  $r^{\text{th}}$  round requires  $2^{(R-r)}$  adder circuits. If we begin with  $s$  qubits per Fourier state going into the first round and double the number of qubits in each subsequent round, then the total number of Toffoli gates in the entire distillation protocol is

$$C_{\text{Tof}} = \sum_{r=1}^R 2^{(R-r)} (2^{(r+1)} s - 4) = 2^{(R+1)} R s - 2^{(R+2)} + 4. \quad (18)$$

Since  $R$  scales as  $O(\log n)$ , then  $C_{\text{Tof}}$  scales as  $O(n \log n)$ . We use  $s = 5$  since the output of the first round is accurate to about 5 bits of precision. Eqn. (17) gives an exact expression for  $R$ . Fig. 4 plots the number of Toffoli gates required for distillation up to  $n = 100$  bits of precision, which is the most precision one could imagine needing for a quantum algorithm. For example, 10 bits of precision is more than sufficient for 4096-bit Shor's algorithm [34]. In this case the approximately 100 Toffoli gates needed to distill  $|\gamma^{(1)}\rangle$  are negligible in comparison to the rest of the algorithm [6]. We emphasize that even if multiple copies of a Fourier state are required, the distillation need only be performed once. Fourier states of size  $n$  qubits can be cloned using a single adder circuit requiring  $(2n - 4)$  Toffoli gates.

The last round of distillation uses  $n + O(1)$  qubits for each input register, which may be less than  $2^{(R-1)} s$ . The additive constant appears because one might need to distill to  $(n+1)$  or  $(n+2)$  qubits in the final output, compensating for errors introduced by truncating the size of Fourier states in earlier rounds. Each round of symmetric distillation doubles the number of qubits per Fourier state, but the number of Fourier states is reduced by half. Consequently, the circuit width of the protocol is at most  $2n + O(1)$  qubits, because the final round uses two registers of size  $n + O(1)$  qubits.




Figure 4: Expected number of Toffoli gates consumed in the distillation of an  $n$ -qubit Fourier state. Uncertainty in number of Toffoli gates is a result of distillation success being probabilistic in each round. No other non-Clifford gates are required, and the Toffoli gates dominate the cost of the adder circuits within distillation.

## 4 Resource comparison for approximating rotation gates

One of the main reasons to distill Fourier states is that they facilitate fault-tolerant phase rotations with phase kickback [14], so we determine the resource costs of this method and compare it to alternatives. Approximating an arbitrary phase rotation with error probability  $\sin^2(\pi/2^{n-1})$  requires an  $n$ -qubit, distilled  $|\gamma^{(1)}\rangle$  Fourier state, including the residual error probability from distillation. As shown previously, preparing such a state requires  $O(n \log n)$  Toffoli gates, but this initialization need only be performed once. Each phase rotation uses an addition circuit, which requires at most  $2n - 4$  Toffoli gates [26]; however, one of the addends is a known value, so we can simplify the adder and “short-circuit” half of the Toffoli gates, replacing them with Clifford gates. In this special case, a single-qubit phase rotation with a precision of  $\pi/2^{n-1}$  radians, which is  $(n - 1)$  bits of precision, requires just  $(n - 2)$  Toffoli gates,  $n$  qubits for the Fourier states, and  $(n - 1)$  ancilla qubits for the internal carry operations of the adder. Forming controlled-rotation gates is simple as well. Each additional control input to the multi-qubit gate requires one more Toffoli gate and one more ancilla qubit.

The best results for approximating single-qubit rotations with sequences of  $T$  gates and Clifford gates each require about 3  $T$  gates per bit of precision [12, 20]. These methods require at most 2 ancilla qubits instead of  $(2n - 1)$  for phase kickback, but the total resource costs from non-Clifford gates is higher. Reference [20] estimates the number of  $T$  gates is  $C_T = 3.21 \log_2(1/\epsilon_F) - 6.93$ ; gate-error  $\epsilon_F$  is defined in Ref. [12] as

$$\epsilon_F = \sqrt{1 - \frac{1}{2} \left| \text{tr}(U^\dagger \tilde{U}) \right|}, \quad (19)$$

where  $\tilde{U}$  is the fault-tolerant sequence approximating gate  $U$ . In the case of a phase kickback rotation accurate to  $p$  bits, the error would be

$$\epsilon_F = \sqrt{1 - \frac{1}{2} |1 + \exp(i\pi/2^p)|} \approx \frac{1}{\sqrt{8}} \left( \frac{\pi}{2^p} \right), \quad (20)$$

where approximation is correct to at least four significant figures for  $p \geq 6$ . A phase kickback rotation accurate to  $p$  bits has the succinct error expression  $\log_2(1/\epsilon_F) \approx p - 0.15$ .

We give direct comparison of resource costs in the two methods approximating rotation gates. To construct a rotation gate accurate to  $p$  bits, one would require (in terms of non-Clifford gates):

- $(3.21p - 6.45)$   $T$  gates (on average) using an approximation sequence [20],  
or
- $(p - 1)$  Toffoli gates using phase kickback [14].

One could produce a Toffoli gate using 4  $T$  gates [24,25], in which case approximation sequences and phase kickback have similar costs in terms of non-Clifford gates. However, efficient fault-tolerant constructions have been found for the Toffoli gate which further reduce the required resources to about the cost of a single  $T$  gate, or even lower [24]. Consequently, the total cost of non-Clifford gates is expected to be lower when using phase kickback. Given that the cost of either  $T$  or Toffoli is much greater than a Clifford gate or ancilla qubit (by about two orders of magnitude [7,9]), there is considerable evidence that phase kickback is the current lowest-overhead construction for a fault-tolerant phase rotation. Further work is needed to develop explicit expressions for resource costs as a function of hardware parameters and the methods of error correction selected.

## 5 Distilling arbitrary Fourier states

We outline here a procedure for distilling any  $|\gamma^{(k)}\rangle$  Fourier state, leaving the detailed analysis for future work. Arbitrary values of  $k$  are needed for QVR phase kickback [14], which is useful in quantum simulation and some implementations of the linear-systems algorithm [14, 35]. Ref. [14] gives a method for transforming any  $n$ -qubit  $|\gamma^{(k)}\rangle$  with odd  $k$  to any other  $|\gamma^{(k')}\rangle$ , using  $O(n^2)$  gates. Since it requires phase kickback with successively larger addition circuits, the number of Toffoli gates is

$$\sum_{s=3}^{n-1} (s-2) = \frac{(n-3)(n-2)}{2}. \quad (21)$$

That protocol is deterministic and does not require distillation, assuming one already has a Fourier state. The protocol in Sec. 2.2 could distill  $|\gamma^{(1)}\rangle$ , which could then be transformed into any  $|\gamma^{(k)}\rangle$ .

Approximations of any  $|\gamma^{(k)}\rangle$  can also be distilled using the protocol in Section 2.2. To have reasonably good efficiency, the approximate initial states need to have substantial fidelity with the desired pure state, say  $F > 0.5$ . There are at least two possible approaches. One is to start  $|\gamma^{(0)}\rangle = |+\rangle^{\otimes n}$ , then apply QVR for each ‘1’ bit in the binary representation of  $k$ , using a  $|\gamma^{(1)}\rangle$  state truncated to  $O(\log n)$  qubits. Each QVR gate need only be accurate to error  $O(2^{-\log n}) = O(1/n)$ , as there are at most  $n$  such operations, so the aggregate error is of order unity, meaning we can bound it below 0.5. These approximate states are then distilled, but they are all of size  $n$  qubits, so the total number of Toffoli gates is  $O(n^2)$ . Whether this method is more efficient than the deterministic construction is not yet clear. A second way to prepare approximate states is to split the quantum register encoding a desired Fourier state into two registers of roughly equal size, and first prepare these approximately through distillation. This method can be applied recursively until the input states are small, like the  $s = 5$  starting state above. For the protocol in Section 3, many of the intermediate registers are  $|\gamma^{(0)}\rangle$ , which can be constructed with Clifford

gates; this will not always hold for arbitrary  $k$ , so the number of Toffoli gates needed here is higher than the distillation of  $|\gamma^{(1)}\rangle$  and may also be  $O(n^2)$  for this method. Since Fourier states used for QVR phase kickback can be reused, the number of *different* values of  $k$  used by an algorithm will dictate whether seeking optimized state-preparation protocols is a worthwhile endeavor.

## 6 Discussion

Although we only use  $Z$  and  $S$  gates to initialize approximate Fourier states, one could also use smaller-angle magic states  $R_z(\pi/2^x)|+\rangle$  for  $x > 2$ , which would increase success probability and decrease the number of rounds. However, this approach would require distillation of those small-angle magic states or approximation of the small-angle rotations [22]. Adding a low-fidelity  $T$  gate to produce a more accurate initial state might be advantageous, but dramatic improvements using smaller-angle rotations are not expected for typical quantum-computing parameters. Conversely, Fourier states would readily enable the complex-instruction-set computing of Ref. [22], because each Fourier state is the tensor product of the desired small-angle magic states. Fourier states can be cloned with just  $(2n - 4)$  Toffoli gates using QVR phase kickback [14]. As such, phase kickback is a better way to produce these magic states than distilling them individually.

Since recent circuit constructions have substantially lowered the cost of a fault-tolerant Toffoli gate [24, 25], one should consider whether further improvements to  $T$  gates are possible. Reductions in  $T$ -gate cost could make approximation sequences a better choice than phase kickback for approximating gates, and this comparison determines whether our results in Fourier-state distillation are useful. We give two arguments for why substantial lowering of  $T$ -gate costs is unlikely. First, magic-state distillation for  $T$  gates places considerable constraints on the underlying quantum codes. It has been conjectured that a protocol distilling states from error  $\epsilon$  to  $O(\epsilon^d)$  yields an output/input distillation fraction less than  $1/d$  [18], and the only known codes that approach this limit are so complicated as to make implementation impractical [19]. From this we infer that magic-state distillation for  $T$  gates cannot improve substantially unless a completely different method is discovered. Second, error correction for Toffoli gates is much more efficient by comparison, and there is further room for improvement. In the best Toffoli constructions, one uses  $T$  gates (at lower fidelity though perhaps distilled) inside an error-detecting circuit to suppress gate errors; with this approach, logical Toffoli is the only non-Clifford gate visible to the algorithm. If the conjecture above were extended to resource costs for Toffoli gates, then the protocol in Refs. [24, 25] saturates the limit through one round, which is why a Toffoli gate costs about the same resources as a  $T$  gate with current methods; further improvements are expected to reduce the cost of a logical Toffoli gate below that of a logical  $T$  gate.

Recent work suggests “V-basis” rotations may also be efficient at approximating arbitrary rotation gates [36, 37]. These non-Clifford gates can be gener-

ated from a distillation procedure that uses  $T$  gates [36]. Further work is needed to determine the fault-tolerant resource costs of V-basis methods and how these costs compare to those of phase kickback or other methods.

By demonstrating how to distill Fourier states and hence complete the phase kickback protocol, this paper has two broader implications. (1) Quantum computing with only Clifford and Toffoli gates is efficient, and simple constructions for arbitrary gates are known. (2) Quantum algorithms using an instruction set of (Clifford+Toffoli) gates could require fewer resources than those using an instruction set of (Clifford+ $T$ ) gates, so the Toffoli gate deserves serious attention as the non-Clifford operation in fault-tolerant quantum computing.

## References

- [1] John Preskill. Reliable quantum computers. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 454(1969):385–410, 1998.
- [2] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 1 edition, October 2000.
- [3] Bei Zeng, A. Cross, and I.L. Chuang. Transversality versus universality for additive quantum codes. *Information Theory, IEEE Transactions on*, 57(9):6272–6284, sept. 2011.
- [4] Bryan Eastin and Emanuel Knill. Restrictions on transversal encoded quantum gate sets. *Phys. Rev. Lett.*, 102:110502, Mar 2009.
- [5] N. Isailovic, M. Whitney, Y. Patel, and J. Kubitowicz. Running a quantum circuit at the speed of data. In *35th International Symposium on Computer Architecture, 2008 (ISCA '08)*, 2008.
- [6] N. Cody Jones, Rodney Van Meter, Austin G. Fowler, Peter L. McMahon, Jungsang Kim, Thaddeus D. Ladd, and Yoshihisa Yamamoto. Layered Architecture for Quantum Computing. *Phys. Rev. X*, 2:031007, 2012.
- [7] Austin G. Fowler and Simon J. Devitt. A bridge to lower overhead quantum computation. *Preprint arXiv:1209.0510v3*, 2012.
- [8] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86:032324, Sep 2012.
- [9] Austin G. Fowler, Simon J. Devitt, and Cody Jones. Surface code implementation of block code state distillation. *Preprint arXiv:1301.7107v1*, 2013.
- [10] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal clifford gates and noisy ancillas. *Phys. Rev. A*, 71(2):022316, Feb 2005.



- [11] Christopher M. Dawson and Michael A. Nielsen. The Solovay-Kitaev Algorithm. *Quantum Inf. Comput.*, 6:81, 2006.
- [12] Austin G. Fowler, Ashley M. Stephens, and Peter Groszkowski. High-threshold universal quantum computation on the surface code. *Phys. Rev. A*, 80(5):052312, 2009.
- [13] Craig R. Clark, Tzvetan S. Metodi, Samuel D. Gasster, and Kenneth R. Brown. Resource requirements for fault-tolerant quantum simulation: The ground state of the transverse Ising model. *Phys. Rev. A*, 79(6):062314, Jun 2009.
- [14] N Cody Jones, James D Whitfield, Peter L McMahon, Man-Hong Yung, Rodney Van Meter, Alán Aspuru-Guzik, and Yoshihisa Yamamoto. Faster quantum chemistry simulation on fault-tolerant quantum computers. *New Journal of Physics*, 14(11):115023, 2012.
- [15] Adam M. Meier, Bryan Eastin, and Emanuel Knill. Magic-state distillation with the four-qubit code. *Preprint* arXiv:1204.4221v1, 2012.
- [16] Matthew Amy, Dmitri Maslov, Michele Mosca, and Martin Roetteler. A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits. *Preprint* arXiv:1206.0758v2, 2012.
- [17] Peter Selinger. Quantum circuits of T-depth one. *Preprint* arXiv:1210.0974v1, 2012.
- [18] Sergey Bravyi and Jeongwan Haah. Magic-state distillation with low overhead. *Phys. Rev. A*, 86:052329, Nov 2012.
- [19] Cody Jones. Multilevel distillation of magic states for quantum computing. *Preprint* arXiv:1210.3388v1, 2012.
- [20] Vadym Kliuchnikov, Dmitri Maslov, and Michele Mosca. Practical approximation of single-qubit unitaries by single-qubit quantum Clifford and T circuits. *Preprint* arXiv:1212.6964v1, 2012.
- [21] Matthew Amy, Dmitri Maslov, and Michele Mosca. Polynomial-time T-depth Optimization of Clifford+T circuits via Matroid Partitioning. *Preprint* arXiv:1303.2042v1, 2013.
- [22] Andrew J. Landahl and Chris Cesare. Complex instruction set computing architecture for performing accurate quantum  $Z$  rotations with less magic. *Preprint* arXiv:1302.3240v1, 2012.
- [23] Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A. Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Phys. Rev. A*, 52:3457–3467, Nov 1995.

- [24] Cody Jones. Low-overhead constructions for the fault-tolerant Toffoli gate. *Phys. Rev. A*, 87:022328, 2013.
- [25] Bryan Eastin. Distilling one-qubit magic states into Toffoli states. *Preprint* arXiv:1212.4872v1, 2012.
- [26] Steven A. Cuccaro, Thomas G. Draper, Samuel A. Kutin, and David Petrie Moulton. A new quantum ripple-carry addition circuit. *Preprint* arXiv:quant-ph/0410184, February 2004.
- [27] Rodney Van Meter and Kohei M. Itoh. Fast quantum modular exponentiation. *Phys. Rev. A*, 71(5):052320, May 2005.
- [28] Thomas G. Draper, Samuel A. Kutin, Eric M. Rains, and Krysta M. Svore. A logarithmic-depth quantum carry-lookahead adder. 6:351–369, July 2006.
- [29] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca. Quantum algorithms revisited. 454(1969):339–354, 1998.
- [30] Alexei Yu. Kitaev, Alexander H. Shen, and Mikhail N. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, 1 edition, 2002.
- [31] Adriano Barenco, Artur Ekert, Kalle-Antti Suominen, and Päivi Törmä. Approximate quantum fourier transform and decoherence. *Phys. Rev. A*, 54:139–146, Jul 1996.
- [32] Charles H. Bennett, Gilles Brassard, Sandu Popescu, Benjamin Schumacher, John A. Smolin, and William K. Wootters. Purification of noisy entanglement and faithful teleportation via noisy channels. *Phys. Rev. Lett.*, 76:722–725, Jan 1996.
- [33] David Deutsch, Artur Ekert, Richard Jozsa, Chiara Macchiavello, Sandu Popescu, and Anna Sanpera. Quantum privacy amplification and the security of quantum cryptography over noisy channels. *Phys. Rev. Lett.*, 77:2818–2821, Sep 1996.
- [34] Austin G. Fowler and Lloyd C. L. Hollenberg. Scalability of shor’s algorithm with a limited set of rotation gates. *Phys. Rev. A*, 70:032329, Sep 2004.
- [35] B. D. Clader, B. C. Jacobs, and C. R. Sprouse. Quantum algorithm to calculate electromagnetic scattering cross sections. *Preprint* arXiv:1301.2340v3, 2013.
- [36] Guillaume Duclos-Cianci and Krysta M. Svore. A State Distillation Protocol to Implement Arbitrary Single-qubit Rotations. *Preprint* arXiv:1210.1980v1, 2012.
- [37] Alex Bocharov, Yuri Gurevich, and Krysta M. Svore. Efficient Decomposition of Single-Qubit Gates into  $V$  Basis Circuits. *Preprint* arXiv:1303.1411v1, 2013.